arXiv:1512.03976v1 [stat.CO] 12 Dec 2015

# An iterative importance sampler for Bayesian parameter estimation in stochastic models of multicellular clocks

Inés P. Mariño[1], Joaquín Míguez[2] and Alexey Zaikin[3]
[1] Department of Biology and Geology, Physics and Inorganic Chemistry,
Universidad Rey Juan Carlos.
Tulipán s/n, 28933 Móstoles, Madrid (Spain). Email: ines.perez@urjc.es
[2] School of Mathematical Sciences, Queen Mary University of London.
Mile End Rd, E1 4NS London (UK). E-mail: j.miguez@qmul.ac.uk.
[3] Institute for Women's Health and Department of Mathematics,
University College London.
London W1T 7DN, United Kingdom. E-mail: alexey.zaikin@ucl.ac.uk

December 15, 2015

**Abstract**

We investigate a stochastic version of the synthetic multicellular clock model proposed by Garcia-Ojalvo, Elowitz and Strogatz. By introducing dynamical noise in the model and assuming that the partial observations of the system can be contaminated by additive noise, we enable a principled mechanism to represent experimental uncertainties in the synthesis of the multicellular system and pave the way for the design of probabilistic methods for the estimation of any unknowns in the model. Within this setup, we investigate the use of an iterative importance sampling scheme, termed nonlinear population Monte Carlo (NPMC), for the Bayesian estimation of the model parameters. The algorithm yields a stochastic approximation of the posterior probability distribution of the unknown parameters given the available data (partial and possibly noisy observations). We prove a new theoretical result for this algorithm, which indicates that the approximations converge almost surely to the actual distributions, even when the weights in the importance sampling scheme cannot be computed exactly. We also provide a detailed numerical assessment of the stochastic multicellular model and the accuracy of the proposed NPMC algorithm, including a comparison with the popular particle Metropolis-Hastings algorithm of Andrieu *et al.*, 2010, applied to the same model and an approximate Bayesian computation sequential Monte Carlo method introduced by Mariño *et al.*, 2013.

**Keywords:** repressilator; intercellular networks; population Monte Carlo; ...

## 1 Introduction

The field of systems biology is rich in problems that demand sophisticated computational tools for estimation, detection and prediction. With the genomics revolution and rise of systems biology, we are witnessing the development of a rigorous engineering discipline to create, control and programme cellular behaviour [3]. The resulting field, known as synthetic biology, has undergone a dramatic growth

throughout the past decade and is poised to transform biotechnology and medicine. A core issue in synthetic biology is the analysis of networks of interacting biomolecules [6], which carry out many essential functions in living cells. However, the design principles underlying the functioning of such intracellular networks remain poorly understood. It is expected, though, that the ability to design synthetic networks will lead both to the engineering of new cellular behaviours and to an improved understanding of naturally occurring networks. A particular system that has drawn considerable attention is the so-called repressilator [14] which is an oscillating network that periodically induces the synthesis of a green fluorescent protein as a readout of its state in individual cells and can be considered as synthetic biological clock. Mathematical models, consisting of systems of nonlinear differential equations, that describe the dynamics of the original repressilator and subsequent extensions of it have appeared in the literature [14–16,20,21,25,26] and sparked interest from research in physics, engineering and mathematics.

In this paper we investigate the application of a novel method termed nonlinear population Monte Carlo (NPMC) [18] to the estimation of a subset of the static parameters of a modified version of this repressilator model [16, 21, 25]. In particular, we introduce a stochastic version of the chaotic, continuous-time modified repressilator model of [21], which consists of a set of stochastic differential equations (SDEs) driven by Wiener noise process. These equations depend on a number of unknown parameters, which we model as random variables. We convert the system of SDEs into an (approximate) discrete-time state space model using a standard Euler-Maruyama scheme and then consider the problem of computing the posterior probability distribution of the unknown parameters in the model conditional on a sequence of partial observations that consist of noisy measurements of a small subset of the (dynamic) state variables. This setup resembles the scenario considered in [21] but (i) the system in this paper is stochastic, while in [21] only a deterministic model was studied, and (ii) we pose a data-poor problem, with the collected observations being low dimensional (2-dimensional, for a 14-dimensional state space), noisy and sparse in time, whereas in [21] data were assumed available continuously in time and noise-free. The randomness in the proposed model dynamics can potentially account for experimental uncertainties in the synthesis of the biological system.

We tackle the Bayesian estimation of the unknown model parameters, i.e., the approximation of their posterior probability distribution conditional on the available observations. We propose to apply the recently introduced NPMC methodology of [18], which involves an iterative importance sampling scheme where the weights undergo a nonlinear transformation to control their variance and mitigate the well-known degeneracy problem of importance samplers [13, 18]. In [18] it was proved that the estimates produced at each iteration of an NPMC algorithm converge in probability and almost surely (a.s.) as the number of Monte Carlo samples, $M$, increases. Therefore, the weight transformation improves the finite-

sample-size performance of the importance sampling scheme, while preserving asymptotic convergence. In the problem tackled in the present paper, however, the weights cannot be computed exactly, but only approximated using a standard bootstrap filter (see, e.g., [13]), in a way similar to the particle Markov chain Monte Carlo (PMCMC) method [1] and some of the numerical examples in [18]. Based on certain unbiasedness properties of particle filters we rigorously prove that the importance sampler with transformed weights also attains asymptotic convergence when using *approximate* weights, even if the complexity of the bootstrap filters used for the approximations is fixed (and hence they introduce errors that do not vanish asymptotically). The latter feature is known as *exact approximation* in the PMCMC literature. Unlike the theorem in [18], the analysis in this paper yields an explicit almost sure convergence rate of the form $M^{-\frac{1}{2}+\epsilon}$, where $M$ is the number of MC samples and $\epsilon > 0$ is an arbitrarily small constant. We have run computer simulations to illustrate how the parameters of the stochastic modified repressilator can be estimated using the proposed NPMC method with noisy partial observations. We have also compared the relative performance of the NPMC method, the approximate Bayesian computation sequential Monte Carlo (ABC SMC) technique [21, 24] and a particle Metropolis-Hastings (PMH) algorithm. Our results show that the NPMC method is more accurate than the state of the art PMH algorithm or the ABC SMC technique with a similar, or even lighter, computational cost.

The rest of the paper is organised as follows. In Section 2 we introduce a stochastic coupled-repressilator model of intercellular networks and derive a discrete-time version. The NPMC algorithm for this system is outlined in Section 3, and a novel theoretical argument for asymptotic convergence is introduced in Section 4. Some computer simulation results are shown in Section 5 and, finally, Section 6 is devoted to the conclusions.

# 2 Intercellular network model

## 2.1 Modified stochastic repressilator

The standard repressilator is a "genetic clock" built by three genes, where the protein product of each gene represses the expression of another one in a cyclic manner [25]. It produces nearly harmonic oscillations in protein levels. In the original repressilator design, the gene *lacI* expresses protein LacI, which inhibits transcription of the gene *tetR*. The product of the latter, TetR, inhibits transcription of the gene *cI*. Finally, the protein product CI of the gene *cI* inhibits expression of *lacI* and completes the cycle. We can see this mechanism in the left side of Figure 1, where the genes are represented in light blue colour.

Figure 1 represents a modification of the repressilator, introduced in [16], that includes an additional feedback loop involving a small autoinducer (AI) molecule produced by CI that can diffuse through the

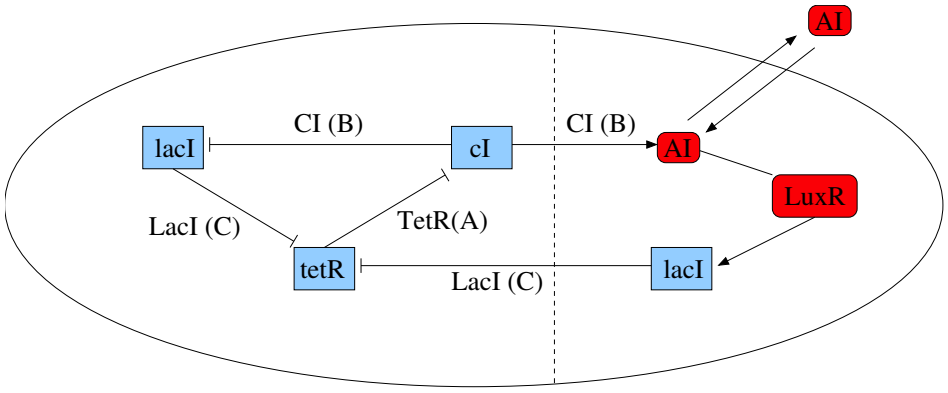


Figure 1: Modified repressilator, where the modification is shown in the right side of the plot. Genes and molecules are represented in light blue and red colours, respectively. The oval depicts the cell membrane.

cell membrane, and the protein LuxR, which responds to the AI by activating the transcription of a second copy of the repressilator gene *lacI*. Placing the gene *cI* under inhibitory control of the repressilator protein TetR leads to a repressive and phase-repulsive coupling that, in turn, generates rich dynamical patterns, including chaotic oscillations [25]. Phase repulsive coupling is common in many biological systems, e.g., in neural activity, in the brain of songbirds or in the respiratory system.

In this paper we study a model consisting of two modified repressilators with identical parameters, driven by Wiener-type noise and coupled by the fast diffusion of the AI across the cell membranes. The resulting mRNA dynamics in continuous time $t \in \mathbb{R}$ is described by a stochastic Hill-type equation with coefficient $\mathsf{m}$, namely

$$\mathsf{d}a_i = -\left(a_i - \frac{\alpha}{1+C_i^{\mathsf{m}}}\right)\mathsf{d}t + \sigma_a a_i \mathsf{d}W_i^a, \tag{1}$$

$$\mathsf{d}b_i = -\left(b_i - \frac{\alpha}{1+A_i^{\mathsf{m}}}\right)\mathsf{d}t + \sigma_b b_i \mathsf{d}W_i^b, \tag{2}$$

$$\mathsf{d}c_i = -\left(c_i - \frac{\alpha}{1+B_i^{\mathsf{m}}} - \frac{\kappa S_i}{1+S_i}\right)\mathsf{d}t + \sigma_c c_i \mathsf{d}W_i^c, \tag{3}$$

where the subscript $i = 1, 2$ specifies the cell; $a_i$, $b_i$, and $c_i$ are time-varying state variables (stochastic processes) representing the concentrations of mRNA molecules transcribed from the genes of *tetR*, *cI*, and *lacI*, respectively; the constant parameter $\alpha$ is the dimensionless transcription rate in the absence

of a repressor; the constant parameter $\kappa$ is the maximum transcription rate of the LuxR promoter; $S_i$ is a state variable representing the concentration of the AI molecule inside cell $i$, and $W_i^a$, $W_i^b$, $W_i^c$, $i = 1, 2$, are independent standard Wiener processes scaled by the constant non-negative factors $\sigma_a$, $\sigma_b$ and $\sigma_c$, respectively. The additional time-varying states $A_i$, $B_i$, and $C_i$, $i = 1, 2$, are stochastic processes representing the concentration of the proteins TetR, CI, and LacI, respectively, whose dynamics obey the SDEs

$$\mathsf{d}A_i = \beta_a(a_i - A_i)\mathsf{d}t + \sigma_A A_i \mathsf{d}W_i^A, \tag{4}$$

$$\mathsf{d}B_i = \beta_b(b_i - B_i)\mathsf{d}t + \sigma_B B_i \mathsf{d}W_i^B, \tag{5}$$

$$\mathsf{d}C_i = \beta_c(c_i - C_i)\mathsf{d}t + \sigma_C C_i \mathsf{d}W_i^C. \tag{6}$$

The equations above show that the dynamics of the proteins is linked to the amount of the responsible mRNA, and the constant parameters $\beta_a$, $\beta_b$ and $\beta_c$ describe the ratio between mRNA and the protein lifetimes (i.e, the inverse degradation rates). Similar to (1)-(3), the dynamics is driven by independent standard Wiener processes $W_i^A$, $W_i^B$ and $W_i^C$, $i = 1, 2$, with constant scale factors $\sigma_A, \sigma_B, \sigma_C \geq 0$. The model is made dimensionless by measuring time in units of the mRNA lifetime (assumed equal for all genes) and the mRNA and protein levels in units of their Michaelis constant. The mRNA concentrations are additionally rescaled by the ratio of their protein degradation and translation rates [16].

The term $\frac{\kappa S_i}{1+S_i}$ on the right-hand side of Eq. (3) represents activated production of *lacI* by the AI molecule, whose concentration inside cell $i$ is denoted by $S_i$. The dynamics of CI and LuxI can be considered identical, given that their production is controlled by the same protein (TetR). Hence, the synthesis of the AI is controlled by the concentration $B_i$ of the protein CI. Taking also into account the intracellular degradation of the AI and its diffusion, the dynamics of $S_i$ is modelled as

$$\mathsf{d}S_i = -\left(k_{s0}S_i - k_{s1}B_i + \eta(S_i - S_e)\right)\mathsf{d}t + \sigma_S S_i \mathsf{d}W_i^S, \tag{7}$$

where $k_{s0}$, $k_{s1}$ and $\eta$ are constant parameters, the latter being a diffusion coefficient that depends on the permeability of the membrane to the AI. The variable $S_e$ is the extracellular concentration of the AI molecule. It is common to apply a quasi-steady-state approximation to the dynamics of $S_e$ [16], which leads to $S_e = Q\overline{S} \equiv Q\frac{1}{N}\sum_{i=1}^{N} S_i$, where $Q = \frac{\delta N}{V_{ext}\left(k_{se}+\frac{\delta N}{V_{ext}}\right)}$, $N = 2$ is the number of cells, $V_{ext}$ is the total extracellular volume, $k_{se}$ is the extracellular AI degradation rate, and $\delta$ is the product of the membrane permeability and the surface area.

This model can produce a range of dynamic regimes. We achieve an underlying chaotic behaviour for this model when the constant parameters are set as [25] $Q = 0.85$, $\mathsf{m} = 2.6$, $\alpha = 216$, $\beta_a = 0.85$, $\beta_b = \beta_c = 0.1$, $\eta = 2$, $\kappa = 25$, $k_{s0} = 1$ and $k_{s1} = 0.01$. We will refer to them as standard values.

## 2.2 Numerical integration and state space model

In order to integrate the 14-dimensional SDE described by Eqs. (1)-(7) numerically, we apply the Euler-Maruyama discretisation with integration step $\mathsf{h} << 1$, that can be explicitly written as

$$
\begin{aligned}
a_{i,m+1} &= a_{i,m} - h\left(a_{i,m} - \frac{\alpha}{1+C_{i,m}^{\mathsf{m}}}\right) + \sigma_a a_{i,m} w_{i,m}^{(1)}, && (8)\\
b_{i,m+1} &= b_{i,m} - h\left(b_i(n) - \frac{\alpha}{1+A_{i,m}^{\mathsf{m}}}\right) + \sigma_b b_{i,m} w_{i,m}^{(2)}, && (9)\\
c_{i,m+1} &= c_{i,m} - h\left(c_{i,m} - \frac{\alpha}{1+B_{i,m}^{\mathsf{m}}} - \frac{\kappa S_{i,m}}{1+S_{i,m}}\right) + \sigma_c c_{i,m} w_{i,m}^{(3)}, && (10)\\
A_{i,m+1} &= A_{i,m} + h\beta_a(a_{i,m} - A_{i,m}) + \sigma_A A_{i,m} w_{i,m}^{(4)}, && (11)\\
B_{i,m+1} &= B_{i,m} + h\beta_b(b_{i,m} - B_{i,m}) + \sigma_B B_{i,m} w_{i,m}^{(5)}, && (12)\\
C_{i,m+1} &= C_{i,m} + h\beta_c(c_{i,m} - C_{i,m}) + \sigma_C C_{i,m} w_{i,m}^{(6)}, && (13)\\
S_{i,m+1} &= S_{i,m} - h\left(k_{s0} S_{i,m} - k_{s1} B_i(n) + \eta(S_{i,m} - S_{e,m})\right) + \sigma_S S_{i,m} w_{i,m}^{(7)}, && (14)
\end{aligned}
$$

where $i = 1, 2$ and $\{w_{i,m}^{(1)}, \ldots, w_{i,m}^{(7)}\}$ are independent Gaussian random variables (r.v.'s) with zero mean and variances $\sigma_a^2, \sigma_b^2, \sigma_c^2, \sigma_A^2, \sigma_B^2, \sigma_C^2$ and $\sigma_S^2$, respectively.

This system described by Eqs. (8)-(14) can be compactly written as the multidimensional difference equation

$$\bar{\mathbf{x}}_m = F_\theta(\bar{\mathbf{x}}_{m-1}, \mathbf{w}_m), \tag{15}$$

where $F_\theta : \mathbb{R}^1 4 \to \mathbb{R}^1 4$ is a function that accounts for both the deterministic and the stochastic part of the model and depends on the parameters $\theta$, $\bar{\mathbf{x}}_m = [\bar{\mathbf{x}}_{1,m}^\top, \bar{\mathbf{x}}_{2,m}^\top]^\top$ is the $14 \times 1$ state of the system at discrete time $m \in \mathbb{Z}$, $\bar{\mathbf{x}}_{i,m}$ are the $7 \times 1$ state vectors associated to the two cells, $i = 1, 2$, and $\mathbf{w}_m = [\mathbf{w}_{1,m}^\top, \mathbf{w}_{2,m}^\top]^\top$ is an independent and identically distributed (i.i.d.) sequence of $14 \times 1$ zero-mean Gaussian vectors. Each $7 \times 1$ subvector $\mathbf{w}_{i,m}$, $i = 1, 2$, has the same distribution and can be written as $\mathbf{w}_{i,m} = [w_{i,m}^{(1)}, \ldots, w_{i,m}^{(7)}]^\top$. In particular, note that, for each cell,

$$\bar{\mathbf{x}}_{i,m} = [a_{i,m}, b_{i,m}, c_{i,m}, A_{i,m}, B_{i,m}, C_{i,m}, S_{i,m}]^\top,$$

with the continuous-time state variables evaluated at time $t = m\mathsf{h}$, e.g., $a_{i,m} = a_i(t = m\mathsf{h})$ As in [21] all constant parameters are assumed known except $\theta = [Q, \mathsf{m}, \alpha, \beta_a]^\top$, which are unknown and modelled as random variables. We assume uniform and independent a priori probability distributions for each one of these parameters, namely $Q \sim \mathcal{U}(0, 1)$, $\mathsf{m} \sim \mathcal{U}(1, 5)$, $\alpha \sim \mathcal{U}(50, 300)$ and $\beta_a \sim \mathcal{U}(0, 1)$. The parameter vector $\theta$, therefore, takes values on the set $\mathsf{S} = (0, 1) \times (1, 5) \times (50, 300) \times (0, 1)$. We denote the conditional (on $\theta$) Markov kernel that determines the state transition from time $m - 1$ to time $m$ as $\bar{\mathcal{K}}_\theta(\mathsf{d}\mathbf{x}_m | \mathbf{x}_{m-1})$.

In particular, for a Borel set $A \subset \mathbb{R}^{14}$, $\bar{\mathcal{K}}_\theta(A|\mathbf{x}_{n-1})$ is the probability of moving from the point $\bar{\mathbf{x}}_{m-1}$ in the state space to some $\bar{\mathbf{x}}_m \in A$.

Partial and noisy observations of the system are collected every $m_0$ discrete time steps, i.e., every $t_o = m_0\mathsf{h}$ continuous time units. Only the variables $a_i$, $i = 1, 2$, are observable, hence the observations have the form

$$\mathbf{y}_n = \begin{bmatrix} a_{1,nm_o} \\ a_{2,nm_o} \end{bmatrix} + \sigma_y \boldsymbol{\epsilon}_n, \quad n = 1, 2, ... \tag{16}$$

where $\boldsymbol{\epsilon}_n$ is a sequence of independent $2 \times 1$ random vectors (with zero mean and identity covariance matrix) and $\sigma_y > 0$ is a known constant parameter.

In order to put the states and the observations on the same time scale, we define the sequence of states $\{\mathbf{x}_n\}_{n\geq 0}$ as $\mathbf{x}_n \triangleq \bar{\mathbf{x}}_{nm_0}$ and introduce the composite Markov kernel

$$\mathcal{K}_\theta(\mathsf{d}\mathbf{x}_n|\mathbf{x}_{n-1}) = \bar{\mathcal{K}}_\theta(\mathsf{d}\mathbf{x}_n|\bar{\mathbf{x}}_{nm_0-1})\bar{\mathcal{K}}_\theta(\mathsf{d}\bar{\mathbf{x}}_{nm_0-1}|\bar{\mathbf{x}}_{nm_0-2}) \cdots \bar{\mathcal{K}}_\theta(\mathsf{d}\bar{\mathbf{x}}_{(n-1)m_0+1}|\mathbf{x}_{n-1}). \tag{17}$$

For a Borel set $A \subset \mathbb{R}^{14}$, $\mathcal{K}_\theta(A|\mathbf{x}_{n-1})$ is the probability of $\mathbf{x}_n = \bar{\mathbf{x}}_{nm_0} \in A$ conditional on $\mathbf{x}_{n-1} = \mathbf{x}_{(n-1)m_0}$ (and the parameter vector $\theta$).

The pair of sequences $\mathbf{x}_n$ and $\mathbf{y}_n$ yield a discrete-time, Markov state space model [27] conditional on the choice of parameters $\theta$. The model is specified by the prior probability distribution of the state $\mathbf{x}_0$, which we denote as $\mathcal{K}_0(\mathsf{d}\mathbf{x}_0)$, the dynamics of the state sequence $\mathbf{x}_n$, which is given by the Markov kernel $\mathcal{K}_\theta(\mathsf{d}\mathbf{x}_n|\mathbf{x}_{n-1})$, and the conditional pdf of the observations $\mathbf{y}_n$ given the states $\mathbf{x}_n$, which we denote as $l_n(\mathbf{y}_n|\mathbf{x}_n)$. We note that, in this model, the latter density is independent of the parameters $\theta$. Also, since $\mathbf{y}_n - [a_{1,nm_0}, a_{2,nm_0}]^\top = \boldsymbol{\epsilon}_n$, the form of the function $l_n(\mathbf{y}_n|\mathbf{x}_n)$ is given by the pdf of the noise term $\boldsymbol{\epsilon}_n$. We often refer to $l_n(\mathbf{y}_n|\mathbf{x}_n)$ as the likelihood of the state $\mathbf{x}_n$.

# 3 Algorithm

In a Bayesian probabilistic setup, the unknown parameters are modelled as a random vector and the aim is to approximate its posterior probability distribution, conditional on the available observations $\mathbf{y} = \{\mathbf{y}_1, \mathbf{y}_1, \ldots, \mathbf{y}_R\}$ for some fixed $R > 0$. We denote the posterior pdf of the parameters as $p(\theta|\mathbf{y})$ and note that it can be readily factored, using Bayes' theorem, as $p(\theta|\mathbf{y}) \propto \ell(\mathbf{y}|\theta)p_0(\theta)$, where $\ell(\mathbf{y}|\theta)$ is proportional to the conditional pdf of the observations $\mathbf{y}$ given the parameters $\theta$ (i.e., the likelihood of $\theta$) and $p_0(\theta)$ is the prior density of $\theta$, which has been chosen to be uniform on the set $\mathsf{S}$, as described in Section 2.2.

The NPMC algorithm of [18] is an iterative importance sampling (IS) scheme that seeks to approximate a target probability distribution, in our case given by the pdf $p(\theta|\mathbf{y})$, using weighted Monte Carlo samples.

This algorithm generates a sequence of proposal pdf's $q_k(\theta)$, $k = 1, \ldots, K$, from which samples can be drawn and importance weights (IWs) can be computed. This sequence of proposals is expected to yield increasingly better approximations of the target as the algorithm converges. The key feature of the NPMC method, which departs from the classical PMC technique of [5], is to compute a set of *transformed* importance weights (TIWs) by applying a nonlinear function to the standard IWs. The aim of this transformation is to mitigate the well known problem of the degeneracy of the IWs (common to many IS methods, see [13, 18]) by controlling the weight variability.

For the case of general state space models, an additional difficulty encountered when trying to estimate the unknown model parameters (denoted $\theta$ in our setup) is that the likelihood $\ell(\mathbf{y}|\theta)$ is intractable. In the last few years, though, it has become a common approach to approximate this likelihood via particle filtering (PF) (see, e.g., [1, 7, 18, 22]). To be specific, we let $\ell^N(\mathbf{y}|\theta)$ stand for the approximation of $\ell(\mathbf{y}|\theta)$ computed using a standard bootstrap filter (BF) [12, 17] with $N$ particles (see Appendix F in [18] for full details). One key feature of this approach, that we exploit for our analysis in Section 4, is that $\ell^N(\mathbf{y}|\theta)$ can be proved to be an unbiased estimator of $\ell(\mathbf{y}|\theta)$ [9, 10].

The NPMC algorithm applied to the model described in Section 2, with $K$ iterations, $M$ Monte Carlo samples per iteration, plain Gaussian proposals $\{q_k\}_{k\geq 1}$, and approximate likelihoods is outlined below.

***Initialisation***. Draw $M$ i.i.d. samples $\theta_0^1, \theta_0^2, \ldots, \theta_0^M$ from the prior pdf $p_0(\theta)$. Then,

1. compute non-normalised IWs $\tilde{\mathsf{w}}_0^i \propto \ell^N(\mathbf{y}|\theta_0^i)$, $i = 1, ..., M$,
2. compute TIWs as $\hat{\mathsf{w}}_0^i = \mathcal{T}_M\left(i, \{\tilde{\mathsf{w}}_0^j\}_{j=1}^M\right)$, where $\mathcal{T}_M : \{1, \ldots, M\} \times \{\tilde{\mathsf{w}}_0^j\}_{j=1}^M \to [0, +\infty)$ is a nonlinear random map (note that the IWs $\tilde{\mathsf{w}}_0^i$ are random),
3. normalise the TIWs, $\mathsf{w}_0^i = \frac{\hat{\mathsf{w}}_0^i}{\sum_{j=1}^M \hat{\mathsf{w}}_0^j}$, $i = 1, ..., M$.

***Iterative step***. For $k = 1, \ldots, K$, take the following steps:

1. Let $q_k(\theta) = N(\theta|\mu_k, \Sigma_k)$ be a multivariate Gaussian pdf with mean vector and covariance matrix obtained, respectively, as

$$\mu_k = \sum_{i=1}^M \mathsf{w}_{k-1}^i \theta_{k-1}^i \quad \text{and} \quad \Sigma_k = \sum_{i=1}^M \mathsf{w}_{k-1}^i \left(\theta_{k-1}^i - \mu_k\right)\left(\theta_{k-1}^i - \mu_k\right)^\top.$$

   Note that the random variates $\theta_{k-1}^i$, $i = 1, ..., M$, are $4 \times 1$ vectors.

2. Draw $\theta_k^i$, $i = 1, ..., M$, i.i.d. samples from $q_k(\theta)$.
3. Compute IWs, $\tilde{\mathsf{w}}_k^i = \frac{\ell^N(\mathbf{y}|\theta_k^i)p_0(\theta_k^i)}{q_k(\theta_k^i)}$, $i = 1, ..., M$.
4. Compute TIWs, $\hat{\mathsf{w}}_k^i = \mathcal{T}_M\left(i, \{\tilde{\mathsf{w}}_k^j\}_{j=1}^M\right)$, $i = 1, ..., M$, using the same nonlinear map as for $k = 0$.
5. Normalise the TIWs, $\mathsf{w}_k^i = \frac{\hat{\mathsf{w}}_k^i}{\sum_{j=1}^M \hat{\mathsf{w}}_k^j}$, $i = 1, ..., M$.

Following [18], the nonlinear map $\mathcal{T}_M$ of choice is a "clipping" transformation. In particular, let $i_1, i_2, ..., i_M$ be a permutation of the indices $1, 2, ..., M$ such that the IWs become ordered, namely $\tilde{\mathsf{w}}_k^{i_1} \geq \tilde{\mathsf{w}}_k^{i_2} \geq \cdots \geq \tilde{\mathsf{w}}_k^{i_M}$. The clipping transformation $\mathcal{T}_M$, with parameter $1 \leq M_c \leq \sqrt{M}$, flattens the $M_c$ largest IWs and makes them equal to the $M_c$-th non-normalised IW, $\tilde{\mathsf{w}}_k^{i_{M_c}}$. Specifically, for each $i = 1, ..., M$, we obtain

$$\hat{\mathsf{w}}_k^j = \mathcal{T}_M\left(j, \{\tilde{\mathsf{w}}_k^l\}_{l=1}^M\right) = \begin{cases} \tilde{\mathsf{w}}_k^{i_{M_c}}, & \text{if } \tilde{\mathsf{w}}_k^j \geq \tilde{\mathsf{w}}_k^{M_c}, \\ \tilde{\mathsf{w}}_k^j, & \text{if } \tilde{\mathsf{w}}_k^j < \tilde{\mathsf{w}}_k^{M_c}, \end{cases} . \tag{18}$$

Other choices of $\mathcal{T}_M$ are possible (e.g., tempering schemes [18]).

Let $p_{\mathbf{y}}(\mathsf{d}\theta) = p(\theta|\mathbf{y})\mathsf{d}\mathbf{y}$ denote the posterior probability measure (conditional on the observed data $\mathbf{y}$) of the parameter vector $\theta$. This measure yields the full probabilistic description of $\theta$ given the available observations. If $p_{\mathbf{y}}(\mathsf{d}\theta)$ is available, then we can compute various types of estimators and assess the associated errors. For example, the posterior-mean estimator is

$$\hat{\theta}_* = \int \mathbb{I}_{\mathsf{S}}(\theta)\theta p_{\mathbf{y}}(\mathsf{d}\theta),$$

where

$$\mathbb{I}_A(z) = \begin{cases} 1, & \text{if } x \in A, \\ 0, & \text{otherwise.} \end{cases}$$

This estimator minimises the mean square error (MSE), which, for an arbitrary estimator $\tilde{\theta}$, can also be written as an integral w.r.t. $p_{\mathbf{y}}(\mathsf{d}\theta)$, namely,

$$\mathrm{MSE}(\hat{\theta}) = \int \mathbb{I}_{\mathsf{S}}(\theta)(\theta - \hat{\theta})^2 p_{\mathbf{y}}(\mathsf{d}\theta).$$

The proposed NPMC algorithm yields a sequence of importance sampling (i.e., weighted Monte Carlo) approximations of $p_{\mathbf{y}}(\mathsf{d}\theta)$. To be specific, at each iteration $k$ we obtain the random probability measure

$$p_{\mathbf{y},k}^M(\mathsf{d}\theta) = \sum_{i=1}^M \mathsf{w}_k^i \delta_{\theta_k^i}(\mathsf{d}\theta),$$

where $\delta_{\theta_k^i}$ denotes the Dirac delta measure centred at $\theta_k^i$. Using $p_{\mathbf{y},k}^M(\mathsf{d}\theta)$ we can approximate any parameter estimator. For instance, $\hat{\theta}_* \approx \hat{\theta}_k^M = \sum_{i=1}^M \mathsf{w}_k^i \theta_k^i$ is the approximation of the posterior mean $\hat{\theta}_*$. The corresponding minimum MSE can also be approximately computed as

$$\mathrm{MSE}(\hat{\theta}_k^M) = \sum_{i=1}^M \mathsf{w}_k^i \|\theta_k^i - \hat{\theta}_k^M\|^2.$$

**Remark 1** *The statement of NPMC algorithm in this section is fairly general. It can be applied, with minor adjustments, to a very broad class of problems involving the Bayesian estimation of unknown parameters in state space stochastic models.*

# 4 Analysis

Consider a single iteration $k$ in the NPMC algorithm, with a fixed importance density $q_k \equiv q$. We refer to the random measure $p^M_{\mathbf{y},k}(\mathsf{d}\theta) = \sum_{i=1}^M \mathsf{w}^i_k \delta_{\theta^i_k}(\mathsf{d}\theta)$ computed via the TIWs $\mathsf{w}^i_k$, $i = 1, .., N$, as a nonlinear importance sampling (NIS) approximation of $p_{\mathbf{y}}(\mathsf{d}\theta)$. Our aim in this section is to assess whether $p^M_{\mathbf{y},k}(\mathsf{d}\theta)$ converges towards the true measure $p_{\mathbf{y}}(\mathsf{d}\mathbf{y})$ or not as $M \to \infty$. To do this, here are two issues that need to be handled and make the analysis more difficult compared to a conventional IS method (that relies on the standard IWs, rather than the TIWs). These issues are:

(i) the distortion in the Monte Carlo approximation due to the clipping of the weights, which can be expected to introduce some bias (compared to the use of standard IWs); and

(ii) the impossibility to compute the IWs, and hence the TIWs, exactly, since the likelihood $\ell(\mathbf{y}|\theta)$ is intractable and we work with a particle approximation $\ell^N(\mathbf{y}|\theta)$ instead.

In [18] it was proved that, when the IWs can be computed exactly, the NIS approximation converges almost surely (a.s.) towards the target probability measure as $M \to \infty$, which accounts for (i) above[1]. The problem of the approximate computation of the weights was partially addressed in [19], for a relatively simple case where the errors in the IWs where assumed deterministic and bounded. However, the estimation problem studied in [19] (parameter estimation for $\alpha$-stable distributions using iid data) did not involve any dynamics and the converge analysis only showed an upper bound for the approximation errors that included a deterministic constant (non-vanishing) term proportional to the IW approximation error.

Here, we show stronger analytical results that ensure the a.s. convergence of the NIS approximation when $M \to \infty$ and can only be estimated as $\ell^N(\mathbf{y}|\theta)$, i.e., using a bootstrap filter with a finite and fixed number of particles $N$. Under assumptions which are standard in the classical IS theory, we prove that integrals of the form $\int f(\theta) p^M_{\mathbf{y},k}(\mathsf{d}\theta)$ converge towards $\int f(\theta) p_{\mathbf{y},k}(\mathsf{d}\theta)$ a.s. as $M \to \infty$ and provide explicit convergence rates.

## 4.1 Notation

Since we focus our attention in the NIS scheme alone, i.e., a single iteration of the proposed algorithm, in the remaining of this section we drop the iteration index $k$. Hence, we assume a fixed importance density $q(\theta)$, from where $M$ independent (yet not necessarily identically distributed) Monte Carlo samples, $\theta^1, \theta^2, \ldots, \theta^M$, are drawn. Since the observations $\mathbf{y}$ are assumed arbitrary but fixed, we drop them from

[1] The analysis of [18] does not provide convergence rates, though.

the likelihood notation and write

$$\ell(\theta) \triangleq \ell(\mathbf{y}|\theta) \quad \text{and} \quad \ell^N(\theta) \triangleq \ell^N(\mathbf{y}|\theta).$$

Then, the non-normalised IWs are approximated as

$$\tilde{\mathsf{w}}^i = g^N(\theta^i) \triangleq \frac{\ell(\theta^i)p_0(\theta^i)}{q(\theta^i)},$$

where we have introduced the weight function $g^N = \ell^N p_0/q$ as a shorthand. This weight function is a random approximation of the deterministic function $g = \ell p_0/q$. The support of $g$ is the same as the support of $q$, $\ell$ and $p_0$, denoted $\mathsf{S} \subseteq \mathbb{R}^4$. We assume that $g(\theta) > 0$ for every $\theta \in \mathsf{S}$ as well (a standard assumption in classical IS). It is also apparent that $p_{\mathbf{y}} \propto gq$, with the proportionality constant independent of $\theta$.

The non-normalised TIWs computed via the clipping function (18) are denoted

$$\hat{\mathsf{w}}^i = [\mathcal{T}^M \circ g^N](\theta^i)),$$

where $\circ$ represents function composition and we omit the index argument of (18) for conciseness (its value is clear from the notation in any case). The normalised TIWs are $\mathsf{w}^i = \frac{\hat{\mathsf{w}}^i}{\sum_{j=1}^M \hat{\mathsf{w}}^j}$, and they are used to compute the approximate measure $p_{\mathbf{y}}^M(\mathsf{d}\theta)$.

## 4.2 Assumptions and a preliminary result

Let the state sequence $\mathbf{x}_n$ take values on some set $\mathcal{X} \subseteq \mathbb{R}^{14}$. We make the following assumptions on the conditional pdf of the observations $\mathbf{y}_n$, $n = 1, 2, \ldots, R$, the prior density of the parameters, $p_0(\theta)$, and the importance function $q(\theta)$.

**Assumption 1** *The observation sequence $\mathbf{y}_{1:R}$ is arbitrary but fixed. The functions $l_n(\mathbf{y}_n|\cdot) : \mathcal{X} \to (0, \infty)$, $n = 1, 2, ..., R$, are uniformly upper bounded and bounded away from zero, i.e., there exist finite and positive constants*

$$\|l\|_\infty = \sup_{n \geq 1, \mathbf{x}_n \in \mathcal{X}} l_n(\mathbf{y}_n|\mathbf{x}_n) < \infty \quad \textit{and} \quad l_{\mathsf{inf}} = \inf_{n \geq 1, \mathbf{x}_n \in \mathcal{X}} l_n(\mathbf{y}_n|\mathbf{x}_n) > 0$$

**Assumption 2** *The ratio of pdf's $\frac{p_0(\theta)}{q(\theta)}$ is upper bounded and bounded away from zero on $\mathsf{S}$, i.e., there exist positive and finite constants*

$$\left\| \frac{p_0}{q} \right\|_\infty = \sup_{\theta \in \mathsf{S}} \left| \frac{p_0(\theta)}{q(\theta)} \right| < \infty \quad \textit{and} \quad \left( \frac{p_0}{q} \right)_{\mathsf{inf}} = \inf_{\theta \in \mathsf{S}} \frac{p_0(\theta)}{q_0(\theta)} > 0.$$

**Remark 2** *If the parameter support set $\mathsf{S}$ is compact, then A.1 and A. 2 hold naturally for most models of practical interest.*

The following lemma plays a key role in the asymptotic convergence analysis of the approximation $p_{\mathbf{y}}^M(\mathsf{d}\theta)$. It states that $\ell^N(\theta)$ is an unbiased estimator of the likelihood $\ell(\theta)$ and enables us to show that the NIS scheme converges when $M \to \infty$, even if the number of particles $N$ in the approximation $\ell^N(\theta)$ remains constant.

**Lemma 1** *If Assumption 1 holds then*

$$\max\{\ell(\theta), \ell^N(\theta)\} \leq \|l\|_\infty^R < \infty, \quad \min\{\ell(\theta), \ell^N(\theta)\} \geq l_{\mathsf{inf}}^R \qquad \textit{and} \quad E\left[\ell^N(\theta)\right] = \ell(\theta)$$

*independently of $N$.*

**Proof.** From the definition of $\ell(\theta)$ in Eq. (39) and its estimator $\ell^N(\theta)$ in Eq. (40), it is clear that both $\ell(\theta) \leq \|l\|_\infty^R$ and $\ell^N(\theta) \leq \|l\|_\infty^R$ when $R$ is the number of available observations. It is also straightforward to show that $\ell(\theta) \geq l_{\mathsf{inf}}^R$ and $\ell^N(\theta) \geq l_{\mathsf{inf}}^R$. The fact that $\ell^N(\theta)$ is unbiased is a consequence of [10, Theorem 7.4.2] (see also [9, Lemma 2] for an alternative proof that does not rely on the Feynmann-Kac framework). ■

## 4.3 Asymptotic convergence

In the sequel we look into the approximation of integrals of the form

$$(f, p_{\mathbf{y}}) \triangleq \int \mathbb{I}_{\mathsf{S}}(\theta) f(\theta) p_{\mathbf{y}}(\mathsf{d}\theta),$$

where $\mathbb{I}_{\mathsf{S}}(\theta)$ is an indicator function (namely, $\mathbb{I}_{\mathsf{S}}(\theta) = 1$ if $\theta \in \mathsf{S}$ and $\mathbb{I}_{\mathsf{S}}(\theta) = 0$ otherwise) and $f$ is a bounded real function in the parameter space $\mathsf{S}$. We use $\|f\|_\infty \triangleq \sup_{\boldsymbol{\theta} \in \mathsf{S}} |f(\boldsymbol{\theta})| < \infty$ to denote the supremum norm of a bounded function while the set of bounded functions on $\mathsf{S}$ is denoted $B(\mathsf{S})$. The approximations of interest are

$$(f, p_{\mathbf{y}}) \approx (f, p_{\mathbf{y}}^M) = \sum_{i=1}^M f(\theta^i) \mathsf{w}^i,$$

for any $f \in B(\mathsf{S})$.

The following Theorem yields an explicit upper bound for the (random) approximation error $|(f, p_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})|$. The bound is proportional to $M^{-\frac{1}{2}}$ and, therefore, it vanishes as $M \to \infty$, independently of the number of particles $N$ used in the approximation of the likelihoods $\ell^N(\theta^i)$.

**Theorem 1** *Assume that A.1 and A.2 hold and $M_c \leq \sqrt{M}$. Then, for every $\epsilon \in \left(0, \frac{1}{2}\right)$ (arbitrarily small) and every $f \in B(\mathsf{S})$ there exists a positive and a.s. finite random variable $V_{f,\epsilon}$, independent of $M$ and $M_c$, such that*

$$|(f, p_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})| \leq \frac{V_{f,\epsilon}}{M^{\frac{1}{2}-\epsilon}}.$$

**Proof.** Recall the intractable weight function $g = \ell p_0/q$ and its random estimator $g^N = \ell^N p_0/q$. The integral of any $f \in B(\mathsf{S})$ w.r.t. the posterior measure $p_{\mathbf{y}}(\mathsf{d}\theta) \propto \ell(\theta)p_0(\theta)\mathsf{d}\theta$ can be written as

$$(f, \pi) = \frac{(fg, q)}{(g, q)}$$

by simply noting that $g(\theta)q(\theta)\mathsf{d}\theta = \ell(\theta)p_0(\theta)\mathsf{d}\theta$. Similarly, we can construct an approximation $\hat{p}^M_{\mathbf{y}}(\mathsf{d}\theta) = \frac{1}{\sum_{j=1}^M g^N(\theta^j)} \sum_{i=1}^M g^N(\theta^i)\delta_{\theta^i}(\mathsf{d}\theta)$ of the posterior $p_{\mathbf{y}}(\mathsf{d}\theta)$ using the IWs *before* the clipping, and then obtain the approximate integral

$$(f, \hat{p}^M_{\mathbf{y}}) = \frac{(fg^N, q^M)}{(g^N, q^M)} \tag{19}$$

where $q^M(\mathsf{d}\theta) = \frac{1}{M}\sum_{i=1}^M \delta_{\theta^i}(\mathsf{d}\theta)$. It is simple to show that

$$(f, \hat{p}^M_{\mathbf{y}}) - (f, p_{\mathbf{y}}) = \frac{(fg^N, q^M) - (fg, q)}{(g, q)} + (f, p^M_{\mathbf{y}})\frac{(g, q) - (g^N, q^M)}{(g, q)} \tag{20}$$

and, since $(g, q) = (\ell, p_0) = \int \mathbb{I}_{\mathsf{S}}(\theta)\ell(\theta)p_0(\theta)\mathsf{d}\theta > 0$ and $(f, \hat{p}^M_{\mathbf{y}}) \leq \|f\|_\infty$, Eq. (20) readily yields

$$|(f, \hat{p}^M_{\mathbf{y}}) - (f, p_{\mathbf{y}})| \leq \frac{1}{(\ell, p_0)}\left|(fg^N, q^M) - (fg, q)\right| + \frac{\|f\|_\infty}{(\ell, p_0)}\left|(g, q) - (g^N, q^M)\right|, \tag{21}$$

and, therefore, the problem of computing a bound for $|(f, \hat{p}^M_{\mathbf{y}}) - (f, p_{\mathbf{y}})|$ reduces to computing bounds for errors of the form $|(bg^N, q^M) - (bg, q)|$, where $b \in B(\mathsf{S})$.

Choose any $b \in B(\mathsf{S})$. A simple triangle inequality yields

$$|(bg^N, q^M) - (bg, q)| \leq |(bg^N, q^M) - (bg, q^M)| + |(bg, q^M) - (bg, q)|. \tag{22}$$

Since $q^M = \frac{1}{M}\sum_{i=1}^M \delta_{\theta^i}$, for the second term on the right hand side of (22) we can write

$$\mathbb{E}\left[|(bg, q^M) - (bg, q)|^p\right] = \mathbb{E}\left[\left|\frac{1}{M}\sum_{i=1}^M Z^i\right|^p\right], \tag{23}$$

where the random variables

$$Z^i = b(\theta^i)g(\theta^i) - (bg, q), \quad i = 1, ..., M,$$

are independent, with zero mean (recall the $\theta^{(i)}$'s are i.i.d. draws from $q$) and bounded, because $b$ is bounded and A.1 and A.2 imply that $g < \|l\|^R_\infty \times \left\|\frac{p_0}{q}\right\|_\infty < \infty$. Therefore, it is an exercise in combinatorics to show that

$$\mathbb{E}\left[\left|\frac{1}{M}\sum_{i=1}^M Z^{(i)}\right|^p\right] \leq \frac{\tilde{c}^p\|l\|^{Rp}_\infty \left\|\frac{p_0}{q}\right\|^p_\infty \|b\|^p_\infty}{M^{\frac{p}{2}}}, \tag{24}$$

where $\tilde{c}$ is a constant independent of $M$ and $q$. Combining (24) with (23) readily yields

$$\|(bg, q^M) - (bg, q)\|_p \leq \frac{\tilde{c}\|l\|^R_\infty \left\|\frac{p_0}{q}\right\|_\infty \|b\|_\infty}{\sqrt{M}}. \tag{25}$$

The inequality (25) implies that there exists an a.s. finite random variable $\tilde{U}_{b,\epsilon} > 0$ such that

$$|(bg, q^M) - (bg, q)| \leq \frac{\tilde{U}_{b,\epsilon}}{M^{\frac{1}{2}-\epsilon}}, \tag{26}$$

where $0 < \epsilon < \frac{1}{2}$ is an arbitrarily small constant independent of $M$ (see [8, Lemma 4.1]).

If we expand the first term on the right hand side of (22) we arrive at

$$\begin{aligned} \left|(bg^N, q^M) - (bg, q^M)\right| &= \left| \frac{1}{M} \sum_{i=1}^{M} b(\theta^i) \left(g^N(\theta^i) - g(\theta^i)\right) \right| \\ &= \left| \frac{1}{M} \sum_{i=1}^{M} Z_N^i \right|, \end{aligned} \tag{27}$$

where the r.v.'s $Z_N^i = \frac{b(\theta^i)p_0(\theta^i)}{q(\theta^i)} \left(\ell^N(\theta^i) - \ell(\theta^i)\right)$, $i = 1, 2, ..., M$, are independent (because the samples $\theta^1, \ldots, \theta^M$ are independent) and zero mean, as a result of Lemma 1. Since they are also bounded, namely $|Z_N^i| \leq \|b\|_\infty \|l\|_\infty^R \left\|\frac{p_0}{q}\right\|_\infty$ as a consequence of A.1 and A.2, it is again an exercise to show that (27) implies

$$E\left[\left|(bg^N, q^M) - (bg, q^M)\right|^p\right] \leq \frac{\bar{c}^p \|l\|_\infty^{Rp} \left\|\frac{p_0}{q}\right\|_\infty^p \|b\|_\infty^p}{M^{\frac{p}{2}}} \tag{28}$$

in the same manner as we obtained the inequality (24). Resorting again to [8, Lemma 4.1], from (28) we deduce that there exists an a.s. finite random variable $\bar{U}_{b,\epsilon} > 0$, independent of $M$, such that

$$|(bg^N, q^M) - (bg, q^M)| \leq \frac{\bar{U}_{b,\epsilon}}{M^{\frac{1}{2}-\epsilon}}, \tag{29}$$

where $0 < \epsilon < \frac{1}{2}$ is an arbitrarily small constant independent of $M$.

Taking together (22), (26) and (29) we arrive at

$$|(bg^N, q^M) - (bg, q)| \leq \frac{U_{b,\epsilon}}{M^{\frac{1}{2}-\epsilon}}, \tag{30}$$

where $U_{b,\epsilon} = \tilde{U}_{b,\epsilon} + \bar{U}_{b,\epsilon} \geq 0$ is an a.s. finite r.v. independent of $M$, and $\epsilon \in \left(0, \frac{1}{2}\right)$ can be chosen to be arbitrarily small.

It is now immediate to combine the inequality (21) with the bound in (30). If we choose $b = f$ in order to control the first term on the right hand side of (21), and $b = 1$ in order to control the second term, we readily find that

$$|(f, \hat{p}_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})| \leq \frac{\tilde{V}_{f,\epsilon}}{M^{\frac{1}{2}-\epsilon}}, \tag{31}$$

where

$$\tilde{V}_{f,\epsilon} = \frac{U_{f,\upsilon} + \|f\|_\infty U_{1,\epsilon}}{(\ell, p_0)} > 0$$

is an a.s. finite random variable independent of $M$, and $\epsilon \in \left(0, \frac{1}{2}\right)$ can be chosen arbitrarily small.

Having found the bound in (31), it is now relatively straightforward to compute the desired error rate for $|(f, p_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})|$. We first apply a triangle inequality to obtain

$$|(f, p_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})| \leq |(f, p_{\mathbf{y}}^M) - (f, \hat{p}_{\mathbf{y}}^M)| + |(f, \hat{p}_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})| \tag{32}$$

and (31) directly yields a bound for the second term on the right hand side of (32). For the first term, we note that

$$(f, p_{\mathbf{y}}^M) = \frac{(f[\varphi^M \circ g^N], q^M)}{(\varphi^M \circ g^N, q^M)}, \tag{33}$$

where $\circ$ denotes composition, hence $(\varphi^M \circ g^N)(\boldsymbol{\theta}) = \varphi^M(g^N(\boldsymbol{\theta}))$. If we now use (33) and the expression for $(f, \hat{p}_{\mathbf{y}}^M)$ in (19) we obtain, by the same argument leading to (21), that

$$\begin{aligned} |(f, p_{\mathbf{y}}^M) - (f, \hat{p}_{\mathbf{y}}^M)| &\leq \frac{|(f[\varphi^M \circ g^N], q^M) - (fg^N, q^M)|}{(\varphi^M \circ g^N, q^M)} + \frac{\|f\|_\infty |(\varphi^M \circ g^N, q^M) - (g^N, q^M)|}{(\varphi^M \circ g^N, q^M)} \\ &\leq l_{\mathsf{inf}}^{-R} \left(\frac{p_0}{q}\right)_{\mathsf{inf}} |(f[\varphi^M \circ g^N], q^M) - (fg^N, q^M)| \\ &\quad + l_{\mathsf{inf}}^{-R} \left(\frac{p_0}{q}\right)_{\mathsf{inf}} \|f\|_\infty |(\varphi^M \circ g^N, q^M) - (g^N, q^M)|, \end{aligned} \tag{34}$$

where the second inequality follows from the definition of the clipping transformation $\varphi^M$ and the bound $g^N \geq l_{\mathsf{inf}}^{-R} \left(\frac{p_0}{q}\right)_{\mathsf{inf}}$. The latter bound is easily obtained by combining A.2 and Lemma 1.

In order to use (34), we look into errors of the form $|(b[\varphi^M \circ g^N], q^M) - (bg^N, q^M)|$ for arbitrary $b \in B(\mathsf{S})$. This turns out relatively straightforward since, from the construction of $\varphi^M$,

$$|(b[\varphi^M \circ g^N], q^M) - (bg^N, q^M)| = \left| \frac{1}{M} \sum_{r=1}^{M_c} b(\theta^{i_r}) \left[ g^N(\theta^{i_{M_c}}) - g^N(\boldsymbol{\theta}^{i_r}) \right] \right| \leq 2\|l\|_\infty^R \left\| \frac{p_0}{q} \right\|_\infty \|b\|_\infty \frac{M_c}{M},$$

where the inequality follows from the bound $g^N \leq \|l\|_\infty^R \left\| \frac{p_0}{q} \right\|_\infty$, which is a straightforward consequence of assumptions A.1 and A.2 and the definition of the estimate $g^N$. We can plug (35) into (34) twice, first choosing $b = f$ and then $b = 1$, in order to control the two terms in the triangle inequality. As a result, we arrive at the *deterministic* bound

$$|(f, p_{\mathbf{y}}^M) - (f, \hat{p}_{\mathbf{y}}^M)| \leq \frac{4\|l\|_\infty^R l_{\mathsf{inf}}^{-R} \left\| \frac{p_0}{q} \right\|_\infty \left(\frac{p_0}{q}\right)_{\mathsf{inf}} \|f\|_\infty M_c}{M} \leq \frac{4\|l\|_\infty^R l_{\mathsf{inf}}^{-R} \left\| \frac{p_0}{q} \right\|_\infty \left(\frac{p_0}{q}\right)_{\mathsf{inf}} \|f\|_\infty}{\sqrt{M}}, \tag{35}$$

where the second inequality follows from the assumption $M_c \leq \sqrt{M}$ in the statement of Theorem 1.

Plugging (35) and (31) into (32) yields the desired error rate,

$$|(f, p_{\mathbf{y}}^M) - (f, p_{\mathbf{y}})| \leq \frac{V_{f,\epsilon}}{M^{\frac{1}{2}-\epsilon}}, \tag{36}$$

where

$$V_{f,\epsilon} = \tilde{V}_{f,\epsilon} + 4\|l\|_\infty^R l_{\mathsf{inf}}^{-R} \left\| \frac{p_0}{q} \right\|_\infty \left(\frac{p_0}{q}\right)_{\mathsf{inf}} \|f\|_\infty$$

is an a.s. finite random variable and $\epsilon < \frac{1}{2}$ is an arbitrarily small constant, both of them independent of $M$. ■

**Remark 3** *Theorem 1 is a general result regarding nonlinear importance sampling. It holds true for any problem involving the approximation of the posterior probability distribution of the unknown parameters of a state space model as long as Assumptions 1 and 2 hold. It is, therefore, not uniquely linked to the repressilator model of interest in this paper.*

# 5 Computer simulations

We have carried out computer simulations to assess both the proposed model, i.e., the stochastic modified repressilator model described by Eqs. (8)-(14), and the NPMC parameter estimation algorithm of Section 3.

For all the simulations presented here, the model parameters are set to their standard values (see Section 2.1) in order to generate synthetic (i.e., simulated) trajectories for the dynamic variables, $\bar{\mathbf{x}}_{i,m}$, $i = 1, 2$ and $m = 0, 1, \ldots$, and sequences of synthetic observationns $\mathbf{y}_n$, $n = 1, 2, \ldots$, according to Eq.(16). This choice of parameters yields an underlying chaotic (and noisy) dynamics of the state variables, as will be shown below. The integration step of the Euler-Mayurama scheme is $\mathsf{h} = 10^{-3}$ time units. When needed, observations are generated every $m_o = 20$ discrete-time steps of model (8)-(14) (equivalently, every $m_o\mathsf{h} = 0.02$ continuous time units). The observational noise $\boldsymbol{\epsilon}_n$ in Eq. (16) is assumed to be zero-mean Gaussian with identity covariance matrix $\mathbf{I}_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, and the standard deviation parameter in the same Eq. (16) is set to $\sigma_y = 1$.

When the parameters $\sigma_a$, $\sigma_b$, $\sigma_c$, $\sigma_A$, $\sigma_B$, $\sigma_C$ and $\sigma_S$, which control the variance of the process noise variables $w_{i,m}^{(j)}$, $j = 1, \ldots, 7$, $i = 1, 2$, are set to zero, we recover the deterministic modified repressilator of [21], which displays chaotic behaviour. For this specific case, we have run a long simulation of the noiseless system (10000 continuous time units) and used the results to obtain phase diagrams. In particular, Fig.2(a) shows the phase diagram for the variable $b_1$ versus $a_1$, while Fig. 2(b) depicts the phase diagram of $a_2$ versus $a_1$. What we observe are two views of the multidimensional chaotic attractor generated by this system. When we add dynamical noise, by setting $\sigma_a^2 = \sigma_b^2 = \sigma_c^2 = \sigma_A^2 = \sigma_B^2 = \sigma_C^2 = \sigma_S^2 = 0.02^2$, we obtain an stochastic dynamical system. However, if we repeat the experiment to generate long trajectories (with the same initial conditions and the same duration) we obtain two similar phase diagrams, as shown in Figs. 2(c) and 2(d). Indeed, the figures simply depict noisy (i.e., slightly perturbed) versions of the original deterministic attractor. This illustrates the fact that the underlying chaotic dynamics is preserved in the stochastic model, which can account for slight perturbations or uncertainties in the physical system

as well.

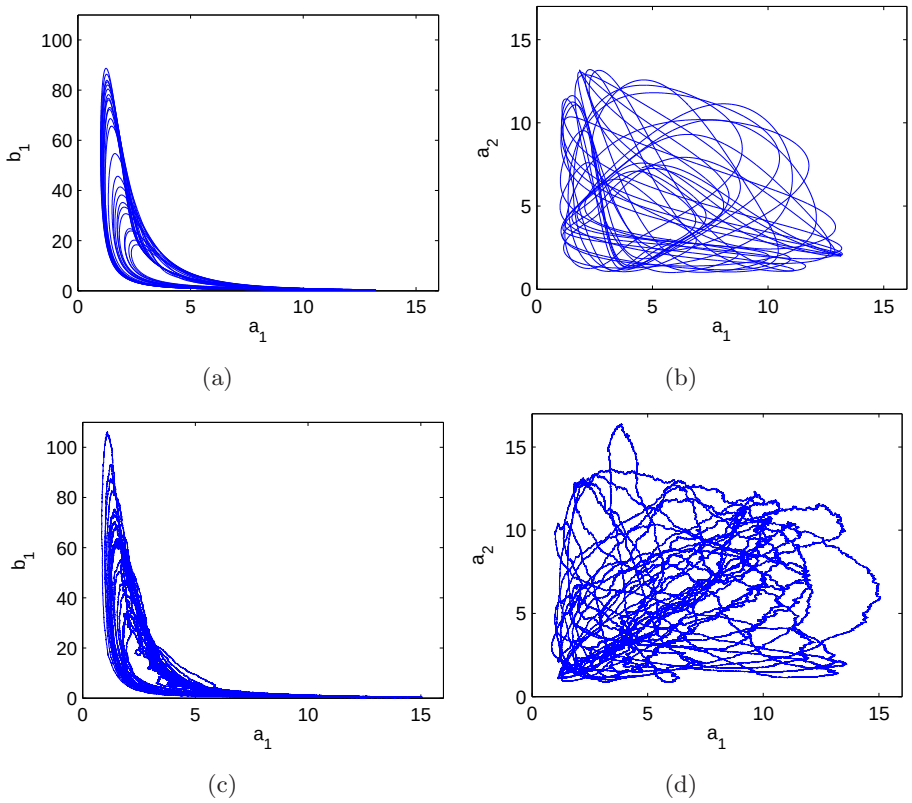


Figure 2: Comparison of 2-dimensional phase space diagrams for the deterministic and the stochastic repressilator models (a) $b_1$ versus $a_1$ and (b) $a_2$ versus $a_1$ for the determinist model; (c) $b_1$ versus $a_1$ and (d) $a_2$ versus $a_1$ for the stochastic model with variance $0.02^2$ for the dynamical noise.

From Figure 2 we also observe that the trajectories of the dynamic variables remain confined within a bounded region. This is relevant because if the state space $\mathcal{X} \subset \mathbb{R}^4$ can be ensured to be compact, then the assumptions A.1 and A.2 on which Theorem 1 in Section 4 relies follow naturally.

In the sequel, we assess the performance of the proposed NPMC method and compare it with some alternative techniques that can be found in the literature. For the subsequent simulation experiments involving the NPMC algorithm, we simulate trajectories of the deterministic modified repressilator model (i.e., we use Eqs. (8)-(14) with $\sigma_a^2 = \sigma_b^2 = \sigma_c^2 = \sigma_A^2 = \sigma_B^2 = \sigma_C^2 = \sigma_S^2 = 0$) with random initial condition[2] and then generate observations over an interval of 80 continuous time units (hence $80/h = 80 \times 10^3$ time steps in the Euler scheme). With observations collected every $m_o = 20$ discrete steps, this yields a sequence of $4 \times 10^3$ 2-dimensional observation vectors contaminated with zero mean Gaussian noise with

[2]The initial condition is generated from a multivariate Gaussian distribution with mean $(a_{1,0}, b_{1,0}, c_{1,0}, A_{1,0}, B_{1,0}, C_{1,0}, S_{1,0}, a_{2,0}, b_{2,0}, c_{2,0}, A_{2,0}, B_{2,0}, C_{2,0}, S_{2,0}) = (4.5, 6, 3, 4.2, 19, 4.3, 0.1, 7.3, 1.5, 3.4, 7, 6.5, 3.6, 0.08)$ and covariant matrix $\sigma_0^2\mathbf{I}$, where $\sigma_0^2 = 0.05^2$.

unit variance. In order to compute the likelihood approximation $\ell^N(\theta)$, which is necessary to obtain IW's and TIW's, we use a bootstrap filter with $N = 100$ particles.

Figure 3 shows the outcome of a sample run of the NPMC algorithm applied to a sequence of $\frac{80}{m_o h}$ observations as described above, with a deterministic trajectory of the state variables and standard parameter values as ground truth. Each plot in Fig. 3 depicts:

- the a priori uniform pdf (solid red line)
- the approximate posterior pdf after one iteration of the NPMC algorithm, and
- the approximate posterior pdf after 15 iterations of the NPMC algorithm,

for an unknown model parameter (hence, for parameters $Q$, $\mathsf{m}$, $\alpha$ and $\beta_a$, from left to right and from top to bottom). Toghether with the densities, the actual parameter value is marked with a vertical dashed-dotted line. The NPMC algorithm was run with $M = 200$ samples per iteration. The approximate pdf's are generated from these samples, $\{\theta_k\}_{i=1}^M$, for iteration $k = 1$ and $k = 15$, using a Gaussian-kernel estimator.

From the plots, we observe that the probability mass of the approximations tends to concentrate around the region where the actual parameter value is located. In this simulation, the true values of $Q$, $\alpha$ and $\beta_a$ are well aligned with the main modes of the approximate pdf's , although this may not necessarily be the case for *all* simulations. Indeed, there is a bias between the actual value of $m$ and the main mode of the approximate density produced by the NPMC method. To explain this "mismatch" let us recall that the approximate statistics generated by the algorithm converge to the true value of these statistics. For example, if we are interested in the posterior expected value of $\theta$, then we compute

$$\hat{\theta}_k^M = \sum_{i=1}^{M} w_k^i \theta_k^i \tag{37}$$

and Theorem 1 ensures that

$$\lim_{M\to\infty} \hat{\theta}_k^M = \int \theta p_{\mathbf{y}}(\theta) d\theta = E[\theta|\mathbf{y}] \quad \text{a.s.} \tag{38}$$

However, depending on the available data (i.e., the dimension of vector $\mathbf{y}$), the posterior mean $E[\theta|\mathbf{y}]$ can be significantly different from the *true* value of $\theta$ used to generate the synthetic data.

A simple way to illustrate this issue is to approximate the likelihood of two different parameter vectors, say $\theta_* = [Q_*, m_*, \alpha_*, \beta_{a*}] = [0.85, 2.6, 216, 0.85]$ the ground truth, and $\theta' = \theta_* + [0, 0, -10, 0]$ a mismatched version, and see that, for a common and fixed observation vector they are approximately the same (actually, $\ell^N(\theta') > \ell^N(\theta_*)$, even if the difference is small). This is shown in Fig. 4, which, for a fixed sequence $\mathbf{y} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n, \ldots\}$, depicts the approximate log-likelihood $\log(\ell^N(\mathbf{y}_{1:n}|\theta_*))$ and

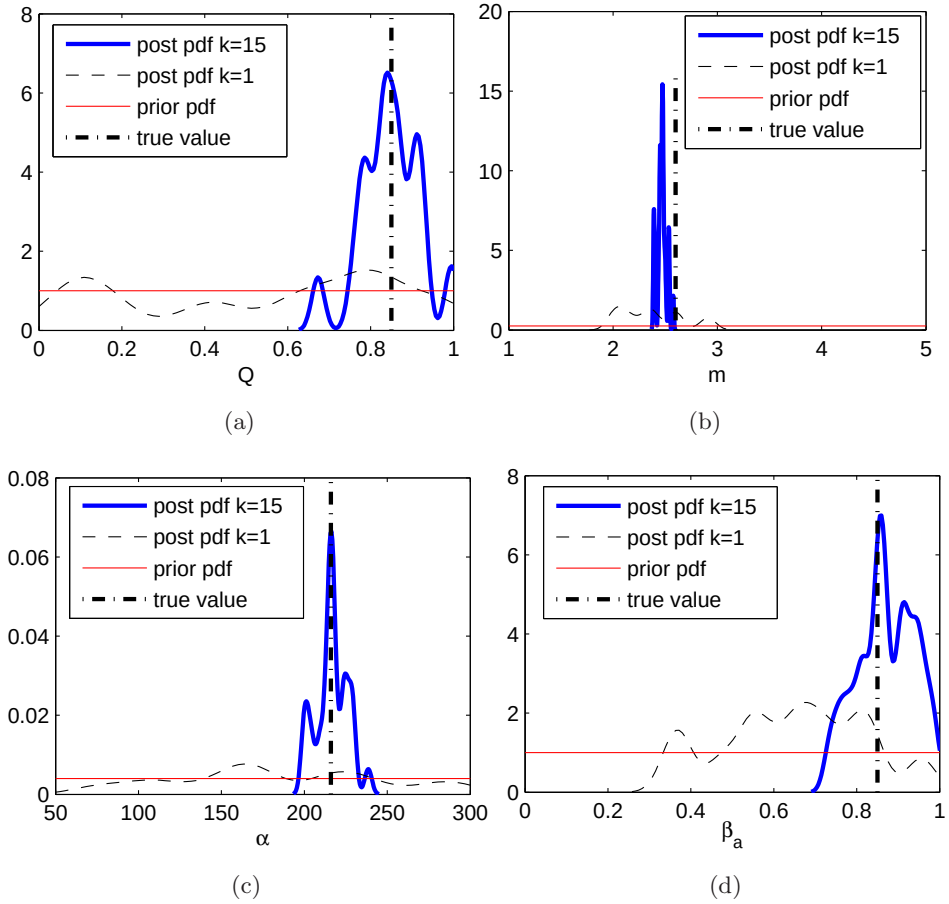


Figure 3: Posterior pdf's computed from the outcome of the NPMC algorithm for $M = 200$ samples per iteration, compared with the true parameter values for the different parameters: (a) posterior density of $Q$, (b) posterior density of $\mathsf{m}$, (c) posterior density of $\alpha$, (d) posterior density of $\beta_a$.

$\log(\ell^N(\mathbf{y}_{1:n}|\theta'))$ versus $n$. The number of particles in the BF is set to $N = 600$ in this case to ensure that we obtain low-variance estimates.

Finally, we have compared the performance of the proposed NPMC method with two other state-of-the-art techniques. The first one is the approximate Bayesian computation (ABC) sequential Monte Carlo (SMC) algorithm of [21] while the second one is a particle Metropolis-Hastings (PMH) [1] method.

ABC SMC is a likelihood-free, Monte Carlo sampling algorithm that relies on a deterministic version of Eqs. (8)-(14) modified to include a coupling term that allows to drive the dynamical system using the observations. The ABC principle involves

- drawing random candidates for the parameter vector, say $\theta^1, \theta^2, \ldots,$
- using these candidates to generate synthetic data, say $\mathbf{y}(\theta^1), \mathbf{y}(\theta^2), \ldots,$ by means of the deterministic

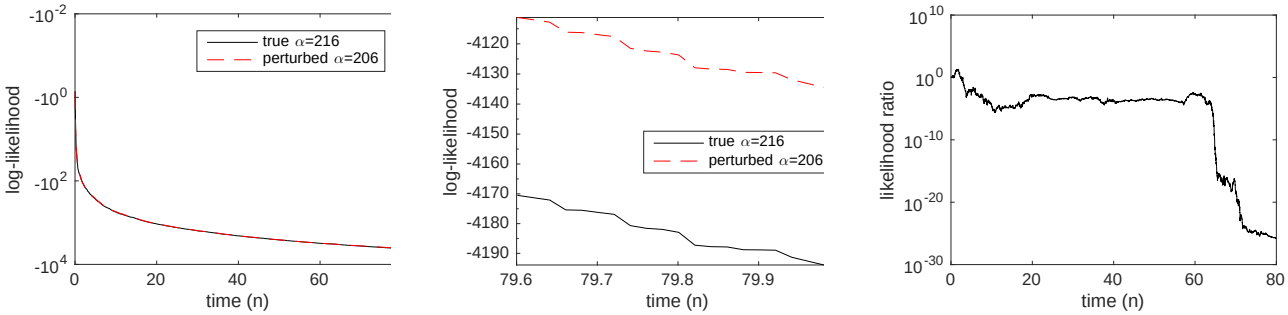


(a) (b) (c)

Figure 4: Comparison of the approximate likelihood of the true parameter vector $\theta_* = [0.85, 2.6, 216, 0.85]$ and perturbed version $\theta' = [0.85, 2.6, \underline{206}, 0.85]$. (a) Approximate log-likelihoods, $\ell(\mathbf{y}_{1:n}|\theta_*)$ and $\ell(\mathbf{y}_{1:n}|\theta')$ versus time $n$. (b) Zoom of plot (a) for a shorter time interval, showing that the perturbed parameter $\theta'$ yields a higher likelihood for the observation sequence $\mathbf{y}$ generated in this computer experiment. (c) The likelihood ratio $\frac{\ell^N(\mathbf{y}_{1:n}|\theta_*)}{\ell^N(\mathbf{y}_{1:n}|\theta')}$ versus time $n$, showing that $\ell^N(\theta') > \ell^N(\theta_*)$.

reppresilator dynamics

- and then comparing the actual and synthetic data using some suitable distance $d$, i.e., evaluating $d(\mathbf{y}, \mathbf{y}(\theta^1)), d(\mathbf{y}, \mathbf{y}(\theta^2)), \ldots$

Samples that yield a small enough distance, typically below a prescribed tolerance, $\epsilon$, are accepted, and those yielding large distances are discarded. Candidates are drawn and tried until a prescribed number of them are accepted. In a SMC setup, this procedure is repeated over several stages with decreasing tolerances $\epsilon_1 > \epsilon_2 > \ldots$ (see [21] for details). Here we have applied the algorithm with tolerances $\epsilon_1 = 3$, $\epsilon_2 = 2.5$, $\epsilon_3 = 2.3$, $\epsilon_4 = 2.2$ and $\epsilon_5 = 2.1$ and the aim of accepting $15 \times 400/5 = 1200$ samples per stage. The maximum number of Monte Carlo draws per stage, however, is set to $1600 \times 10^3$, to prevent the algorithm from getting stuck at any stage due to low acceptance rates.

The PMH algorithm is a Markov chain Monte Carlo (MCMC) technique [23]. It generates candidates with a Gaussian kernel with covariance matrix

$$\mathbf{\Sigma} = \begin{pmatrix} \sigma_Q^2 & 0 & 0 & 0 \\ 0 & \sigma_\alpha^2 & 0 & 0 \\ 0 & 0 & \sigma_{\mathsf{m}}^2 & 0 \\ 0 & 0 & 0 & \sigma_{\beta_a}^2 \end{pmatrix} = \begin{pmatrix} 0.01 & 0 & 0 & 0 \\ 0 & 100 & 0 & 0 \\ 0 & 0 & 0.01 & 0 \\ 0 & 0 & 0 & 0.01 \end{pmatrix}$$

and applies the Metropolis-Hastings rule to accept or reject them. This demands the evaluation of the likelihood $\ell(\mathbf{y}|\theta')$ for each candidate $\theta'$, and therefore it is approximated using a BF, $\ell(\mathbf{y}, \theta') \simeq \ell^N(\mathbf{y}|\theta')$ with $N = 100$, the same as for the proposed NPMC method. We have used the algorithm to generate Markov chains of length $15 \times 400 = 6000$, which is equivalent to the computational cost of the NPMC method with $M = 400$ samples per iteration and $k = 15$ iterations.

We have compared the ABC SMC algorithm, the PMH method and the proposed NPMC scheme with $k = 15$ iterations and $M \in \{50, 200\}$ samples in terms of the empirical normalised mean square error (NMSE) of the posterior-mean estimators of each parameter ($Q$, $\mathsf{m}$, $\alpha$ and $\beta_a$), computed from 45 independent simulation runs.

Figure 5 depicts the outcomes of the comparison. For each parameter, we shows the NMSE (and the NMSE plus one standard deviation) obtained empirically from the simulations, for each algorithm. The NPMC method, even with just $M = 50$ samples, outperforms the PMH and ABC SMC techniques, even if the latter have a significantly larger computational cost. The NPMC algorithm with $M = 200$ samples is still computationally less demanding than the implemented PMH and ABC SMC methods[3] and it attains an improvement in estimation accuracy. This improvement is modest for some parameters, but this is again a consequence of the limited amount of data available for the estimator.

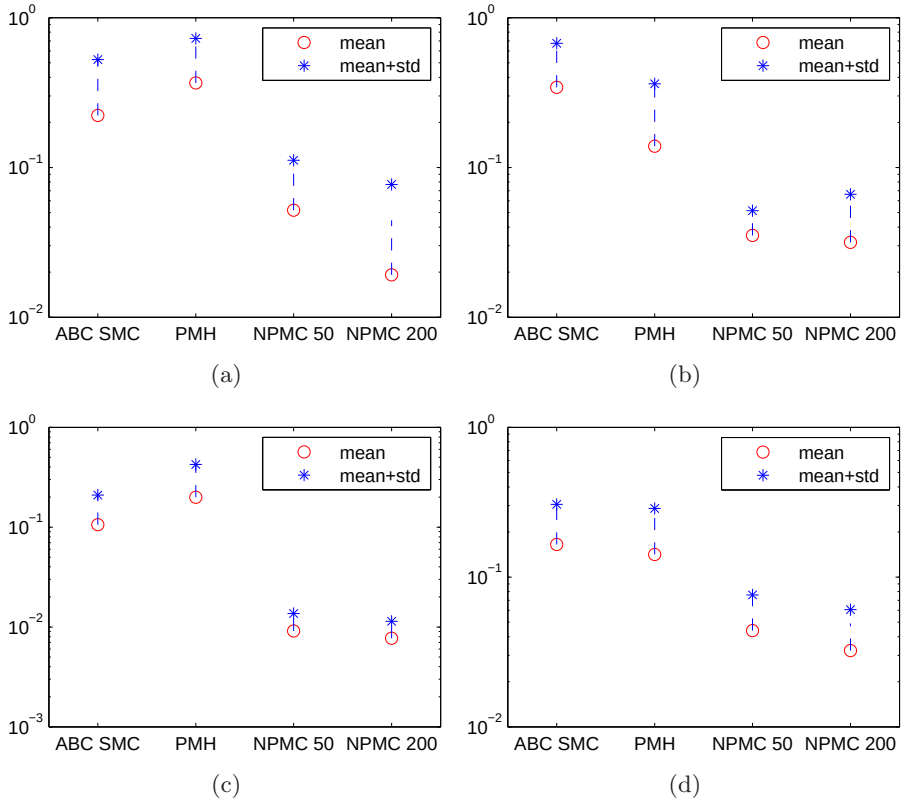


Figure 5: Average values (plus one standard deviation) of the empirical mean square errors over independent 33 simulations for all three methods (ABC SMA, PMH and NPMC with $M = 50$ and $M = 200$ Monte Carlo samples) and each unknown parameter: (a) parameter $Q$, (b) parameter $\mathsf{m}$, (c) parameter $\alpha$, (d) parameter $\beta_a$.

[3]The running time of the ABC SMC algorithm is random because it is based on a rejection sampling scheme. In the average, its computational cost is similar to an NPMC method with 400 samples, although some simulation runs can be much faster.

# 6 Conclusion

We have proposed a stochastic version of the coupled repressilator model of [16] that enables

(a) a mathematically principled manner of describing experimental uncertainties in the synthesis of multicellular clocks, and

(b) the design of probabilistic methods for the estimation of unknown parameters in the model.

In particular, we have introduced an iterative Monte Carlo sampling scheme for the approximation of the posterior probability distributions of the parameters of interest conditional on a set of noisy and partial observations of the system. The technique relies on a methodology termed nonlinear importance sampling, originally introduced in [18]. In this paper, we have applied this methodology to the stochastic repressilator and proved a new theoretical result regarding the convergence of nonlinear importance samplers. The new convergence theorem is stronger than the original results in [18] and holds for a broader class of models (of which the stochastic repressilator system is just an instance). Specifically, we have proved that nonlinear importance samplers can converge asymptotically with optimal Monte Carlo error rates even when the importance weights can only be estimated (and the variance of these estimates is positive and cannot be reduced).

The theoretical analysis has been complemented with an extensive computer simulation study that illustrates the relationship between the deterministic and stochastic repressilator models and demonstrates the efficiency and accuracy of the proposed estimation algorithm compared to other computational approaches of similar complexity that can be found in the literature.

# Acknowledgments

This research has been partially supported by the Spanish Ministry of Economy and Competitiveness (projects TEC2012-38883-C02-01 COMPREHENSION and FIS2013-40653-P), the Spanish Ministry of Education, Culture and Sport (mobility award PRX15/00378), the Office of Naval Research (ONR) Global (Grant Award no. N62909-15-1-2011), the regional Government of Madrid (program CASI-CAM-CM S2013/ICE-2845) and the Cancer Research UK and the Eve Appeal Gynaecological Cancer Research Fund (grant ref. A12677) supported by the National Institute for Health Research (NIHR) University College London Hospitals (UCLH) Biomedical Research Centre.

# A The bootstrap filter

A Markov state-space model consists of two sequences of r.v.'s, $\{\mathbf{x}_n\}_{n\geq 0}$ and $\{\mathbf{y}_n\}_{n\geq 1}$. The first sequence, $\{\mathbf{x}_n\}$, is termed the system state. We assume it takes values on some space $\mathcal{X} \subseteq \mathbb{R}^{d_x}$, hence $\mathbf{x}_n$ is a random $d_x \times 1$ vector. The state dynamics are described by a prior probability measure $\mathcal{K}_0(\mathsf{d}\mathbf{x}_0)$ and a sequence of Markov kernels $\mathcal{K}_{n,\theta}(\mathsf{d}\mathbf{x}_n|\mathbf{x}_{n-1})$ that depend on a parameter vector $\theta \in \mathbb{R}^{d_\theta}$. In the case of the modified stochastic repressilator model, the parameter vector is[4] $\theta = (Q, m, \beta_a, \alpha)$ and the Markov kernel $\mathcal{K}_{n,\theta}$ is given by Eq. (17).

The state $\mathbf{x}_n$ cannot be observed directly. Instead, some partial noisy observations $\mathbf{y}_n$ are collected. We assume that the observations are conditionally independent given the system states and the parameter vector $\theta$, with a conditional pdf (with respect to the Lebesgue measure) $l_{n,\theta}(\mathbf{y}_n|\mathbf{x}_n) > 0$, which depends on the parameters $\theta$ as well. For the stochastic repressilator model, the observations are given by Eq. (16), hence $l_{n,\theta}(\mathbf{y}_n|\mathbf{x}_n) = l_n(\mathbf{y}_n|\mathbf{x}_n)$ is independent of the parameter vector $\theta$ in the case of the repressilator model of interest in this paper.

The bootstrap filter (BF) [12, 17] is a recursive Monte Carlo algorithm for the approximation of the sequence of posterior probability measures $\pi_{n,\theta}(\mathsf{d}\mathbf{x}_n)$, $n = 1, 2, ...$, where

- for a *given* (i.e., fixed, even if arbitrary) sequence of observations $\mathbf{y}_{1:n} = \{\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_n\}$,
- and a Borel set $A \subset \mathcal{X}$,

$\pi_{n,\theta}(A)$ is the probability of the even $\mathbf{x}_n \in A$ conditional on the observations $\mathbf{y}_{1:n}$ and the parameter values given by $\theta$. The BF with $N$ *particles* (i.e., Monte Carlo samples) can be briefly outlined as follows.

1. ***Initialisation.*** Draw $N$ samples $\mathbf{x}_0^1, \ldots, \mathbf{x}_0^N$ from the prior distribution $\mathcal{K}(\mathsf{d}\mathbf{x}_0)$. The particle approximation of $\pi_{0,\theta}(\mathsf{d}\mathbf{x}_0) \equiv \mathcal{K}_0(\mathsf{d}\mathbf{x}_0)$ is

$$\pi_{0,\theta}^N(\mathsf{d}\mathbf{x}_0) = \frac{1}{N}\sum_{i=1}^N \delta_{\mathbf{x}_0^i}(\mathsf{d}\mathbf{x}_0),$$

   where $\delta_{\mathbf{x}_0^i}$ denotes the Dirac (unit) delta measure centred at $\mathbf{x}_0^i \in \mathcal{X}$.

2. ***Recursive step.*** Given the approximation $\pi_{n-1,\theta}^N(\mathsf{d}\mathbf{x}_{n-1}) = \frac{1}{N}\sum_{i=1}^N \delta_{\mathbf{x}_{n-1}^i}(\mathsf{d}\mathbf{x}_{n-1})$, take the following steps:

   (a) Randomly propagate each particle using the Markov kernel in the model, i.e., draw $\tilde{\mathbf{x}}_n^i$ from $\mathcal{K}_{n,\theta}(\mathsf{d}\mathbf{x}_n|\mathbf{x}_{n-1}^i)$, $i = 1, ..., N$.

   (b) Compute IWs, $\tilde{u}_n^i = l_{n,\theta}(\mathbf{y}_n|\mathbf{x}_n^i)$, for $i = 1, ..., N$, and

[4] We are interested on the parameters to be estimated alone. Known parameters are implicitly included in the model.

(c) normalise them as

$$u_n^i = \frac{\tilde{u}_n^i}{\sum_{j=1}^N \tilde{u}_n^j}, \quad i = 1, ..., N.$$

(d) Resample: draw $N$ times independently from the discrete distribution $\tilde{\pi}_{n,\theta}^N(\mathsf{d}\mathbf{x}_n) = \sum_{i=1}^N u_n^i \delta_{\tilde{\mathbf{x}}_n^i}(\mathsf{d}\mathbf{x}_n)$ and denote the resulting samples as $\{\mathbf{x}_n^i\}_{i=1}^N$. Construct the to *unweighted* approximation $\pi_{n,\theta}^N(\mathsf{d}\mathbf{x}_n) = \frac{1}{N}\sum_{i=1}^N \delta_{\mathbf{x}_n^i}(\mathsf{d}\mathbf{x}_n)$.

The resampling step (d) above can be implemented in a number of different ways (see, e.g., [2, 11] or [4] for a brief survey of methods). Here, for simplicity, we have adopted a scheme which is often referred to as multinomial resampling [11, 13] but most asymptotic convergence results hold true for several other schemes as well [2, 10]. The random measure $\pi_{n,\theta}^N$ can be used to approximate integrals w.r.t. the true posterior measure $\pi_{n,\theta}$. To be specifc, it can be shown under mild assumptions on the state space model [2,10] that $\|(f, \pi_{n,\theta}^N) - (f, \pi_{n,\theta})\|_p \leq \frac{C}{\sqrt{N}}$ for every $p \geq 1$, where $f(\mathbf{x}_n)$ is a real bounded function of the state $\mathbf{x}_n$, $\|Z\|_p = (E[Z^p])^{\frac{1}{p}}$ is the $L_p$ norm of the r.v. $Z$, $(f, \pi) := \int f(\mathbf{x})\pi(\mathbf{x})$ denotes the integral of function $f$ w.r.t. the measure $\pi$ and $C$ is a finite constant (independent of $N$).

The algorithm also produces an approximation of the predictive probability measure of $\mathbf{x}_n$ conditional on the observations $\mathbf{y}_{1:n-1}$. We denote the actual predictive measure as $\xi_{n,\theta}(\mathsf{d}\mathbf{x}_n)$ and its $N$-particle approximation as

$$\xi_{n,\theta}^N(\mathsf{d}\mathbf{x}_n) = \frac{1}{N}\sum_{i=1}^N \delta_{\tilde{\mathbf{x}}_n^i}(\mathsf{d}\mathbf{x}_n).$$

If we write $\mathbf{y} = \mathbf{y}_{1:n}$, it turns out that the likelihood of the parameter vector $\theta$, namely $\ell(\mathbf{y}|\theta)$ can actually be expressed in terms of integrals w.r.t. to the sequence of predictive distributions $\xi_{n,\theta}$. To be specific,

$$\ell(\mathbf{y}|\theta) = \prod_{k=1}^t (l_{k,\theta}(\mathbf{y}_k|\cdot), \xi_{k,\theta}) \tag{39}$$

and, therefore, the bootstrap filter yields the straightforward estimator

$$\ell^N(\mathbf{y}|\theta) = \prod_{k=1}^t (l_{k,\theta}(\mathbf{y}_k|\cdot), \xi_{k,\theta}^N) = \frac{1}{N^t}\prod_{k=1}^t \sum_{i=1}^N l_{k,\theta}(\mathbf{y}_k|\tilde{\mathbf{x}}_k^i), \tag{40}$$

which can be proved to be unbiased under very mild assumptions [10].

# References


[1] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society B*, 72:269–342, 2010.

[2] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*. Springer, 2008.

[3] D Ewen Cameron, Caleb J Bashor, and James J Collins. A brief history of synthetic biology. *Nature Reviews Microbiology*, 12(5):381–390, 2014.

[4] O. Cappé, S. J. Godsill, and E. Moulines. An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924, 2007.

[5] O. Cappé, A. Gullin, J. M. Marin, and C. P. Robert. Population monte carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, 2004.

[6] Nathalie Chabrier-Rivier, Marc Chiaverini, Vincent Danos, François Fages, and Vincent Schächter. Modeling and querying biomolecular interaction networks. *Theoretical Computer Science*, 325(1):25–44, 2004.

[7] N. Chopin, P. E. Jacob, and O. Papaspiliopoulos. SMC2: an efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2012.

[8] D. Crisan and J. Miguez. Particle-kernel estimation of the filter density in state-space models. *Bernoulli*, 20(4):1879–1929, 2014.

[9] D. Crisan, J. Miguez, and G. Ríos. A simple scheme for the parallelisation of particle filters and its application to the tracking of complex stochastic systems. *arXiv*, arXiv:1407.8071v2 [stat.CO], 2015.

[10] P. Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, 2004.

[11] R. Douc, O. Cappé, and E. Moulines. Comparison of resampling schemes for particle filtering. In *Proceedings of the $4^{th}$ International Symposium on Image and Signal Processing and Analysis*, pages 64–69, September 2005.

[12] A. Doucet, N. de Freitas, and N. Gordon. An introduction to sequential Monte Carlo methods. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*, chapter 1, pages 4–14. Springer, 2001.

[13] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo Sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.

[14] Michael B Elowitz and Stanislas Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403(6767):335–338, 2000.

[15] R David Evans and Luis A Ricardez-Sandoval. Multi-scenario modelling of uncertainty in stochastic chemical systems. *Journal of Computational Physics*, 273:374–392, 2014.

[16] J. Garcia-Ojalvo, M. B. Elowitz, and S. H. Strogatz. Modeling a synthetic multicellular clock: repressilators coupled by quorum sensing. *Proceedings of the National Academy of Sciences of the United States of America*, 101(30):10955–10960, 2004.

[17] N. Gordon, D. Salmond, and A. F. M. Smith. Novel approach to nonlinear and non-Gaussian Bayesian state estimation. *IEE Proceedings-F*, 140(2):107–113, 1993.

[18] E. Koblents and J. Míguez. A population monte carlo scheme with transformed weights and its application to stochastic kinetic models. *Statistics and Computing*, 25(2):407–425, 2014.

[19] E. Koblents, J. Miguez, M. A. Rodriguez, and A. M. Schmidt. A nonlinear population Monte Carlo scheme for the Bayesian estimation of parameters of $\alpha$-stable distributions. *Computational Statistics and Data Analysis*, 95:57–74, March 2016.

[20] A Koseska, E Ullner, E Volkov, J Kurths, and J Garcia-Ojalvo. Cooperative differentiation through clustering in multicellular populations. *Journal of theoretical biology*, 263(2):189–202, 2010.

[21] I. P. Mariño, E. Ullner, and A. Zaikin. Parameter estimation methods for chaotic intercellular networks. *PLoS ONE*, 8(11):e79892, November 2013.

[22] Vasileios Maroulas and Panos Stinis. Improved particle filters for multi-target tracking. *Journal of Computational Physics*, 231(2):602–611, 2012.

[23] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2004.

[24] Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael PH Stumpf. Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187–202, 2009.

[25] E. Ullner, A. Zaikin, E. I. Volkov, and J. García-Ojalvo. Multistability and clustering in a population of synthetic genetic oscillators via phase-repulsive cell-to-cell communication. *Physical Review Letters*, 99(14):148103, 2007.

[26] Ekkehard Ullner, Aneta Koseska, Jürgen Kurths, Evgenii Volkov, Holger Kantz, and Jordi García-Ojalvo. Multistability of synthetic genetic networks with repressive cell-to-cell communication. *Physical Review E*, 78(3):031904, 2008.

[27] M. West and J. Harrison. *Bayesian Forecasting, 2nd ed.* Springer-Verlag, New York, 1996.